Radiel Health - Uncertainty Aware Prediction for Hemodynamic Parameters

Nicholas Jiang
University of Waterloo
nicholas.jiang@uwaterloo.ca

Rishabh Sharma
University of Waterloo
r342shar@uwaterloo.ca

Ahash Ganeshamoorthy
Griffith University
ahash.ganeshamoorthy@griffithuni.edu.au

***Abstract*—We present an uncertainty-aware neural surrogate for predicting hemodynamic parameters on vascular geometries. Our architecture combines Graph Attention Networks (GAT) for local geometric encoding with General Powerful Scalable (GPS) layers for global context, and a Bayesian Neural Network (BNN) task head for epistemic uncertainty quantification. On bifurcation geometries with varying Reynolds numbers and angles, the model achieves $R^2 = 0.962$ for Wall Shear Stress (WSS) magnitude with a 12% relative error, while providing calibrated 95% credible intervals. This addresses a core challenge in deploying machine learning surrogates for clinical decision support: knowing when not to trust the model.**

***Index Terms*—Hemodynamics, Machine Learning, Uncertainty Quantification, Computational Fluid Dynamics**

I. Introduction

A. Background

a) *Computational Fluid Dynamics for Hemodynamics:* Cardiovascular disease is the second leading cause of death in Canada [1]. This includes pathologies such as Coronary Artery Disease where atherosclerotic plaque deposits in the arterial lumen. Sudden plaque rupture can lead to sudden ischemia and subsequent death [2]. Computational Fluid Dynamics (CFD) simulates fluid flow by numerically solving the Navier-Stokes equations. In cardiovascular medicine, CFD enables the non-invasive analysis of blood flow patterns from patient imaging data. This holds significant potential in improving diagnostics and device design for cardiovascular disease.

b) *Clinical Relevance of Wall Shear Stress:* Wall Shear Stress (WSS) quantifies the frictional force blood exerts on vessel walls. Abnormal WSS is implicated in atherosclerosis, where low WSS regions correlate with plaque formation [3], aneurysm rupture risk, where high WSS gradients indicate vulnerable vessel walls, and arteriovenous fistula maturation, where WSS patterns predict surgical outcomes for dialysis patients. Fast, accurate WSS prediction would enable real-time surgical planning.

c) *ML for CFD:* Neural learning methods for Computational Fluid Dynamics (CFD) have become popular in recent years for their potential speedup over traditional CFD solvers [4]. The literature has been progressing towards increasingly general ML surrogates that preserve accuracy and increase inference speed [5], [6]. In medicine, traditional CFD is not widely adopted because of its long run time and computational inefficiencies. In particular, it requires solving non-linear partial differential equations (PDEs) simultaneously and repeatedly over complex geometries and requires expert experience to set-up and operate [7]. A recent survey by Chiastra et al., 2023 reported that 89.0% of interventional cardiologists would be open to utilizing CFD to identify lesions that may lead to future major cardiovascular events. In fact, 85.9% of respondents would consider using it in their daily practice. However, the study highlighted the necessity for fast, low-cost hemodynamic calculation models to realize the predictive power of CFD. As such, ML for CFD models can bridge this gap between in-silico planning and clinical practice, fulfil the interest that clinicians have exhibited. This motivates ML surrogates that trade exactness for clinically acceptable accuracy at interactive speeds, allowing for up to a 90x inference speed increase compared to traditional CFD [8].

B. Problem Formulation

Our model focuses on steady-state flow regimes, a standard approximation for hemodynamic modelling. This allows us to provide clinically relevant characterizations of cardiovascular behavior. For example, we use Wall Shear Stress (WSS) in a steady-state flow to approximate Time-Averaged Wall Shear Stress (TAWSS). This simplifies our output flow field so it is no longer time dependent while maintaining clinical utility. We aim to approximate the solution operator Φ of the steady-state Navier-Stokes equations, which maps from a domain geometry and flow parameters to a flow field. Formally, given a point cloud $\Omega \subset \mathbb{R}^n$ representing the geometry and flow conditions ξ, we approximate the parametric mapping:

$$\Phi : (\Omega, \xi) \to \left(u_\xi : \Omega \to \mathbb{R}^n\right) \tag{1}$$

where $u_{\xi(\vec{x})} = \mathrm{WSS}(\vec{x})$ denotes the Wall Shear Stress vector at point $\vec{x} \in \Omega$ under flow conditions ξ. With our neural model, we aim to learn weights θ such that $\Phi_\theta \approx \Phi$ under the mean-squared error (MSE):

$$\frac{1}{|\mathcal{D}|} \sum_{(\Omega,\xi)\in\mathcal{D}} \left(\frac{1}{|\Omega|} \sum_{\vec{x}\in\Omega} \left\| \Phi_{\theta(\Omega,\xi)}(\vec{x}) - \Phi(\Omega,\xi)(\vec{x}) \right\|_2^2 \right) \tag{2}$$

where $\mathcal{D} = \left\{ \left(\Omega_i, \xi_j\right) \right\}$ is the dataset of (point cloud, flow condition) pairs. Note that we obtain Φ, the target parametric mapping, via CFD simulations from OpenFOAM and ANSYS Fluent. We opted for a purely data-driven surrogate over physics-constrained architectures (e.g., PINNs) to maximize

inference efficiency. While physics-informed models typically necessitate the evaluation of the governing PDEs across the entire volumetric domain, our approach learns the direct mapping onto the boundary manifold. By bypassing the interior flow field, we aim to achieve a significant reduction in degrees of freedom.

This would enable near-real-time clinical insights into surface-bound metrics like WSS. However, this efficiency introduces a reliance on in-distribution data. If a user specifies a Reynolds number (e.g., $\mathrm{Re} = 10,000$) that far exceeds the training range ($\mathrm{Re} \in [100, 3200]$), or inferences on an entirely unseen geometry, the model will be prone to out-of-distribution (OOD) errors.

II. Methodology

A. Model Architecture

To mitigate the risk of unknown OOD error, we integrated a Bayesian Neural Network (BNN) into the task head. Unlike deterministic models, the BNN captures epistemic uncertainty by generating a predictive distribution rather than a single point estimate [9]. We quantify this by performing $N = 100$ stochastic forward passes for each point on the mesh to calculate a 95% Credible Interval, defined by the 0.025 and 0.975 quantiles. For OOD data, we expect this interval to widen significantly, providing a warning to the clinician.

To provide the necessary spatial context for WSS prediction, our architecture employs a Geometry Encoder composed of GINEConv layers [10], This encoder generates a latent feature vector for each node $x \in \Omega$, effectively embedding the local topological environment into a high-dimensional representation. This design choice is motivated by the physical observation that hemodynamic fields like WSS exhibit local continuity on the vascular manifold. By utilizing isotropic aggregation, the model dynamically learns to combine features from adjacent nodes, allowing sensitivity to local curvature and flow-directing features.

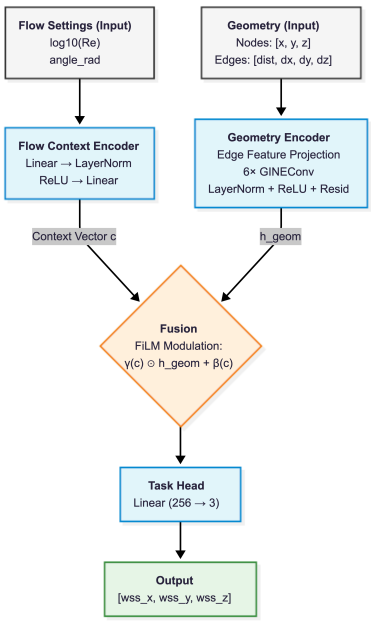

Finally, we input the expanded feature vector representation of each node generated by the Geometry Encoder along with the coordinates (x, y, z) to the task-head where we use BNNs to sample the predictive distribution and return both a 95% credible interval and the arithmetic mean for the WSS vector at every point on the mesh boundary.

B. Evaluation

a) *Predictive Accuracy:*

To evaluate our model's predictive accuracy, we applied it to predicting WSS on variants of a bifurcation mesh where we vary the Reynolds number (Re) between 100 and 2100, the angle of the bifurcation (30, 45, 60 degrees), and the size of the child vessels (500mm, 750mm, 1000mm). The size of the parent vessel stayed the same (1000mm). This gave us 189 graphs to be split across supervised learning and testing. We defined a train/val/test split of 70/15/15, evaluating on 27 held-out cases stratified by Reynolds number and bifurcation angle.

b) *Qualitative Visualization:*

We will visualize the magnitude of the model's WSS predictions and compare it to the ground truth. Earlier, we argued that out model's BNN component allows clinicians to visualize where the model is uncertain in it's predictions in an effort to avoid 'silent' OOD errors. We will also visualize the width of the credible intervals on predicting magnitude of WSS on the geometry and discuss the results.

c) *Generalization:*

In order to see how well the model performs as the test data gets further from the training, we will also set up a smaller experiment with a problem that is simpler to perform perturbations to. The 2D Lid Driven Cavity is a fundamental problem in fluid dynamics and very cheap to run as compared to the 3D bifurcation. It involves a box-like cavity where the top lid moves, creating motion within the domain for the model to predict. We have 11 different sizes of cavity, where there are three domain sizes (in meters) we trained our model on, 1x1, 2x1, and 1x2. All other cases are within the 1mx1m and 1mx2m domains, with a step size of 0.1 meters. We will also define the distance from the training set as the average Euclidean distance from a given point to points in set $S = \left\{ \binom{1}{1}, \binom{1}{2}, \binom{2}{1} \right\}$

(i.e. the aspect ratios we used to trained the network). For a given domain size $\vec{p} \in \mathbb{R}^2$, the distance is:

$$d(\vec{p}) = \frac{1}{|S|} \sum_{\vec{s} \in S} \|\vec{s} - \vec{p}\|_2 \tag{3}$$

So, that means that, from the dataset we have, the 1.9 meter by 1 meter cavity is the farthest from the dataset in our collection of testing points.

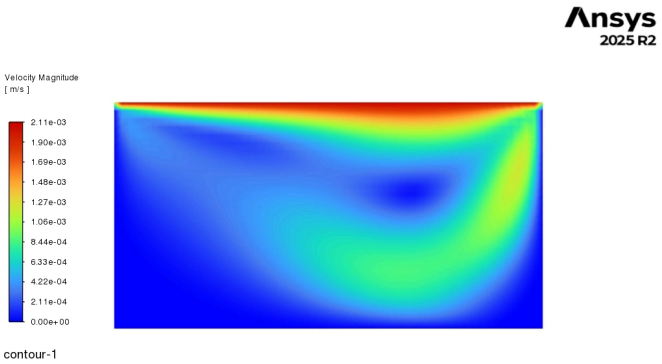

Fig. 1. Results from Ansys Fluent simulations of the 1.9 meter by 1 meter Lid Driven Cavity, the most 'abnormal' out of the dataset we are evaluating with.

We will use the same Reynold's numbers as described above, from 100 to 2100 (steps of 100), bringing us to a total of 189 cases. We have a similar 70/15/15 train/test/val split, evaluating on 28 held-out cases from an equal portion of each cavity shape.

III. Results

A. Predictive Accuracy

Our model is robust across different Reynolds numbers, only performing worse on very small Re < 250. The weaker performance on WSS_Y ($R^2 = 0.503$) reflects a fundamental signal-to-noise challenge: along the direction of flow, $|\text{WSS}_Y|$ and $|\text{WSS}_Z|$ is near-zero, so even minor absolute errors produce poor R^2 scores. Ultimately, the flow along the Z-axis is the only one with clinical relevance as the other flow directions provide no real information about the characteristics of the problem at hand.

TABLE I
Model performance on test set across WSS components.

Component	R^2	NRMSE	Median Rel. Err.
Magnitude	0.962	0.196	12.9%
WSS_Z	0.970	0.175	13.5%
WSS_X	0.955	0.213	45.0%
WSS_Y	0.503	0.705	N/A

B. Uncertainty Quantification

We compare the results of visualizing the credible interval widths on an in-distribution geometry compared to an OOD geometry. Note, all of the bifurcations below have an angle of 45 degrees between the children vessels, each child vessel is 500mm in length and

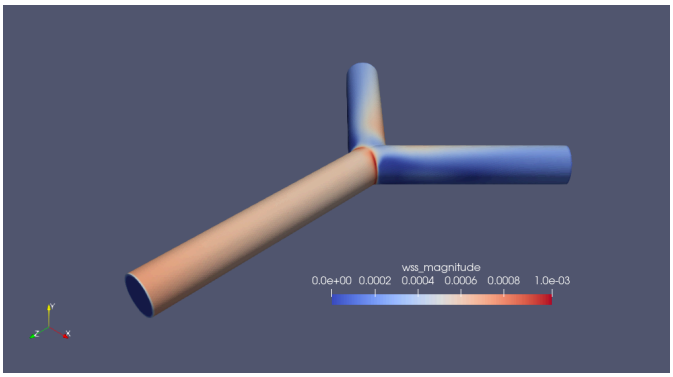

Fig. 2. Predicted WSS magnitude

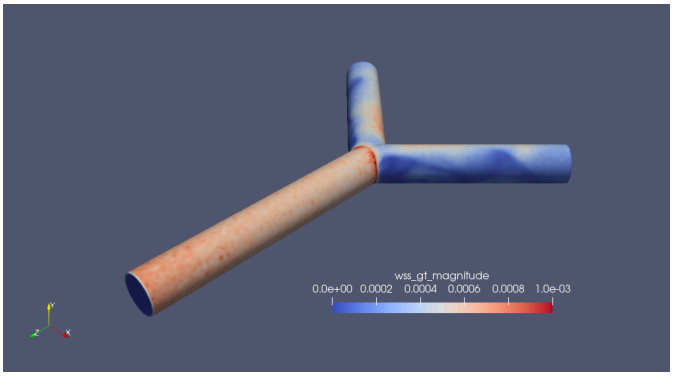

Fig. 3. Ground Truth WSS magnitude

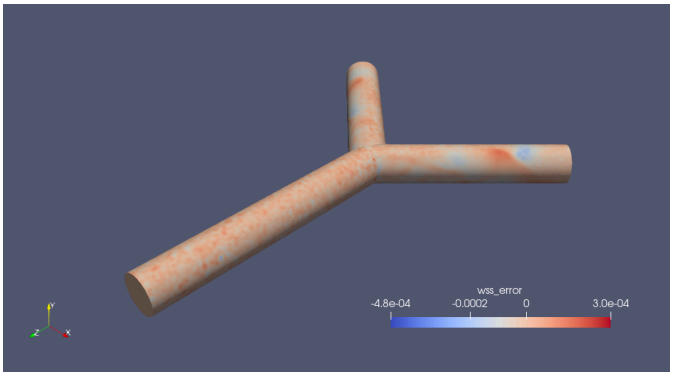

Fig. 4. Absolute difference between prediction and ground truth values

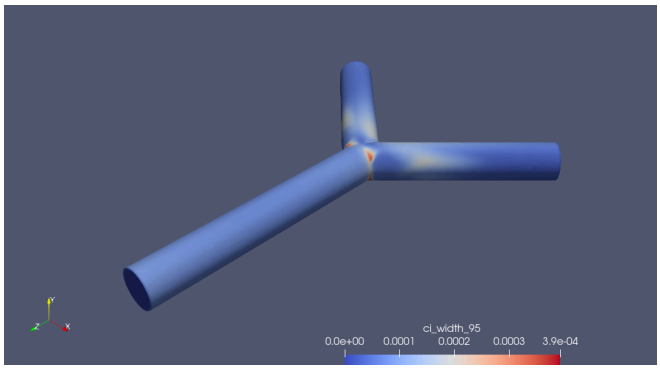

Fig. 5. Width of the 95% credible interval of model predictions for each point on the bifurcation geometry. Red regions indicate higher model uncertainty.

We see that the width of the credible interval gives a good estimate of the absolute difference between the prediction and ground truth values for this case of $\mathrm{Re} = 1600$ and 45° angle bifurcation. The absolute error is distributed close to uniformly on the mesh and the width of the 95% credible intervals are also relatively evenly distributed and are narrow except at the apex of bifurcation.

C. Generalizability

Our model generalizes well on unseen cavity shapes across all Reynold's numbers. The overall evaluation results across all 8 evaluation cavity sizes are below, showcasing a very high variance and a median error between 3-6% all all aspect ratios.

TABLE II
MODEL PERFORMANCE ON UNSEEN GEOMETRIES ACROSS RE OF 100 TO 2100.

Component	R^2	RMSE (Pa)	MAE (Pa)	Median Error
WSS_X	0.995	4.81×10^{-8}	1.99×10^{-8}	3.8%
WSS_Y	0.990	3.77×10^{-8}	1.65×10^{-8}	6.4%
WSS Magnitude	0.998	3.26×10^{-8}	1.74×10^{-8}	3.1%

As expected, the RMSE goes up as the cavity size gets farther from the test set, following the same linear pattern across all Reynold's numbers.

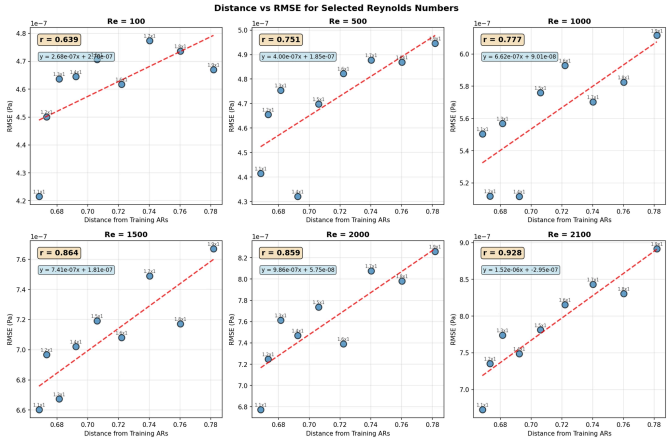

Fig. 6. RMSE predicted across geometries further from the training set

IV. CONCLUSION

We presented an uncertainty-aware neural surrogate for hemodynamic parameter prediction, combining a GINEConv-based geometry encoder and a Bayesian task head for epistemic uncertainty quantification. On bifurcation geometries, the model achieves strong predictive accuracy for WSS magnitude ($R^2 = 0.962$) while providing calibrated credible intervals that flag unreliable predictions. Our model also demonstrated it's ability to generalize well on unseen structures. This architecture addresses a core challenge in deploying ML surrogates for clinical decision support: knowing when not to trust the model. Future work will extend this framework to patient-specific vascular geometries derived from CT imaging and evaluate uncertainty calibration on true out-of-distribution anatomies.

REFERENCES

[1] A. Auverach and others, "The Burden of Cardiovascular Disease," *Circulation*, 2024, doi: 10.1161/CIRCULATIONAHA.124.070628.

[2] C. Chiastra *et al.*, "Computational fluid dynamics as supporting technology for coronary artery disease diagnosis and treatment: an international survey," *Frontiers in Cardiovascular Medicine*, vol. 10, p. 1216796, 2023, doi: 10.3389/fcvm.2023.1216796.

[3] H. Samady *et al.*, "Coronary artery wall shear stress is associated with progression and transformation of atherosclerotic plaque and arterial remodeling in patients with coronary artery disease," *Circulation*, vol. 124, no. 7, pp. 779–788, 2011.

[4] Y. Luo, Y. Chen, and Z. Zhang, "CFDBench: A Large-Scale Benchmark for Machine Learning Methods in Fluid Dynamics," no. arXiv:2310.05963. arXiv, Feb. 2024. doi: 10.48550/arXiv.2310.05963.

[5] R. Zhang, Q. Meng, H. Wan, Y. Liu, Z.-M. Ma, and H. Sun, "OmniFluids: Physics Pre-trained Modeling of Fluid Dynamics." [Online]. Available: https://arxiv.org/abs/2506.10862

[6] H. Zhou, Y. Ma, H. Wu, H. Wang, and M. Long, "Unisolver: PDE-Conditional Transformers Towards Universal Neural PDE Solvers." [Online]. Available: https://arxiv.org/abs/2405.17527

[7] P. D. Morris *et al.*, "Computational fluid dynamics modelling in cardiovascular medicine," *Heart*, vol. 102, no. 1, pp. 18–28, 2016, doi: 10.1136/heartjnl-2015-308044.

[8] R. Gharleghi, A. Sowmya, and S. Beier, "Transient wall shear stress estimation in coronary bifurcations using convolutional neural networks," *Computer Methods and Programs in Biomedicine*, vol. 225, p. 107013, 2022, doi: https://doi.org/10.1016/j.cmpb.2022.107013.

[9] S. Lee, H. Kim, and J. Lee, "Graddiv: Adversarial robustness of randomized neural networks via gradient diversity regularization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

[10] W. Hu *et al.*, "Strategies for Pre-training Graph Neural Networks." [Online]. Available: https://arxiv.org/abs/1905.12265